

WHAT IS CLAIMED IS:

1. A node for use in a multi-node computer system, the node comprising:

5 a plurality of active devices;

an interface configured to send and receive coherency messages on an inter-node network coupling nodes in the multi-node computer system;

10 an address network configured to communicate address packets between the active devices and the interface; and

a data network configured to communicate data packets between the active devices and the interface;

15 wherein the active device includes a promise array configured to store a promise identifying a data packet to be conveyed to a device in response to a pending local transaction involving a coherency unit for which the active device has an ownership responsibility;

20 wherein the active device is configured to store promises in the promise array in response to receiving address packets from other ones of the plurality of active devices and from the interface.

25 2. The node of claim 1, wherein the active device includes a cache subsystem and an interface controller coupled to the cache subsystem, wherein the interface controller is configured to ensure that the active device has at most one outstanding local transaction for a given coherency unit.

3. The node of claim 2, wherein the promise array includes storage for at least one promise associated with each outstanding local transaction.

4. The node of claim 2, wherein for each outstanding local transaction, the promise array includes storage for one promise for each interface included in the node and for one promise for each other one of the plurality of active devices.

5. The node of claim 1, wherein in response to receiving the data packet as part of the pending local transaction, the active device is configured to transition an access right associated with the coherency unit.

6. The node of claim 5, wherein the active device is configured to gain a write access right to the coherency unit in response to receiving the data packet, and wherein in response to the promise, the active device is configured to send the data packet to the device, wherein the active device transitions the write access right to a shared access right in response to sending the data packet to the device.

7. The node of claim 1, wherein the address network includes an interface request virtual network configured to communicate address packets sent by the interface and a non-interface request virtual network configured to communicate address packets sent by devices other than the interface

8. The node of claim 7, wherein the active device is configured to assert flow control on the non-interface request virtual network without asserting flow control on the interface request virtual network.

9. The node of claim 7, wherein the active device is configured to assert flow control on the non-interface request virtual network in response to storing a threshold number of promises in the promise array.

10. The node of claim 7, wherein the address network also includes a multicast virtual network, a response virtual network, and a broadcast virtual network.

11. The node of claim 1, wherein the active device is configured to store the promise
5 in the promise array in response to receiving a write-stream packet sent by the device on the address network, wherein the device is another active device, and wherein the promise identifies an acknowledgment packet to be sent to the device.

12. The node of claim 1, wherein the active device is configured to store the promise
10 in the promise array in response to a read-to-own packet sent by the device on the address network, wherein the device is another active device, wherein the active device is configured to lose the ownership responsibility for the coherency unit upon receipt of the read-to-own packet.

15 13. The node of claim 1, wherein the active device is configured to store the promise in the promise array in response to a proxy-read-to-own-modified packet sent by the interface on the address network, wherein the device is the interface, and wherein the active device is configured to lose the ownership responsibility for the coherency unit upon receipt of the proxy-read-to-own-modified packet.

20

14. A system, comprising:

25 a plurality of nodes, wherein each of the plurality of nodes includes: an active device, an interface configured to send and receive coherency messages on an inter-node network coupling the plurality of nodes; an address network configured to communicate address packets between the active device and the interface; and a data network configured to communicate data packets between the active device and the interface;

wherein the active device includes a promise array configured to store a promise identifying a data packet to be conveyed to a device in response to a pending local transaction involving a coherency unit for which the active device has an ownership responsibility;

5

wherein the active device is configured to store promises in the promise array in response to receiving address packets from the interface and other active devices in a same node of the plurality of nodes as the active device.

10 15. The system of claim 14, wherein the active device includes a cache subsystem and an interface controller coupled to the cache subsystem, wherein the interface controller is configured to ensure that the active device has at most one outstanding local transaction for a given coherency unit.

15 16. The system of claim 15, wherein for each outstanding local transaction, the promise array includes storage for one promise for each interface included in the same node and for one promise for each one of a plurality of active devices in the same node.

17. The system of claim 14, wherein in response to receiving the data packet as part of
20 the pending local transaction, the active device is configured to transition an access right associated with the coherency unit.

18. The system of claim 14, wherein the address network includes an interface request virtual network configured to communicate address packets sent by the interface and a
25 non-interface request virtual network configured to communicate address packets sent by devices other than the interface

19. The system of claim 18, wherein the active device is configured to assert flow control on the non-interface request virtual network without asserting flow control on the
30 interface request virtual network.

20. The system of claim 18, wherein the active device is configured to assert flow control on the non-interface request virtual network in response to storing a threshold number of promises in the promise array.

5

21. The system of claim 14, wherein the active device is configured to store the promise in the promise array in response to a proxy packet sent by the interface on the address network, wherein the device is the interface.

10 22. A method for use in a multi-node computer system, wherein the multi-node computer system includes a plurality of nodes coupled by an inter-node network, the method comprising:

15 an interface included in a node of the plurality of nodes receiving a communication specifying a coherency unit on the inter-node network from another one of the plurality of nodes;

20 in response to receiving the communication, the interface sending an address packet specifying the coherency unit on an address network coupling the interface to an active device in the node;

the active device storing a promise in a promise array in response to having an ownership responsibility for the coherency unit and receiving the address packet from the interface;

25

in response to said storing the promise and receiving a data packet as part of an outstanding local transaction involving the coherency unit, the active device sending another data packet to the interface.

23. The method of claim 22, wherein the active device includes a cache subsystem and an interface controller coupled to the cache subsystem, wherein the interface controller ensures that the active device has at most one outstanding local transaction for a given coherency unit.

5

24. The method of claim 22, wherein for each outstanding local transaction, the promise array includes storage for one promise for each interface included in the node and for one promise for each one of a plurality of other active devices included in the node.

10

25. The method of claim 22, further comprising the active device transitioning an access right associated with the coherency unit in response to said receiving the data packet as part of the pending local transaction.

15

26. The method of claim 22, wherein said the interface sending the address packet comprises the interface sending the address packet on an interface request virtual network, wherein active devices included in the node send address packets on a non-interface request virtual network.

20

27. The method of claim 26, further comprising the active device asserting flow control on the non-interface request virtual network without asserting flow control on the interface request virtual network.

25

28. The method of claim 22, further comprising the active device asserting flow control on the non-interface request virtual network in response to storing a threshold number of promises in the promise array.